

AD-A059 639

DECISIONS AND DESIGNS INC MCLEAN VA

F/G 5/10

RECONCILING INCOHERENT JUDGMENTS (RIJ) - TOWARD PRINCIPLES OF P--ETC(U)

JUL 78 R V BROWN, D V LINDLEY

N00014-75-C-0426

UNCLASSIFIED

TR-78-8-72

NL


| OF |
AD
A059639



END
DATE
FILMED

12-78

DDC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

LEVEL

2

TECHNICAL REPORT TR 78-8-72

AD A059639


RECONCILING INCOHERENT JUDGMENTS (RIJ)- TOWARD PRINCIPLES OF PERSONAL RATIONALITY

DECISIONS AND DESIGNS INCORPORATED

Rex V. Brown
Dennis V. Lindley

July 1978

DDC FILE COPY



This document has been approved
for public release and sale; its
distribution is unlimited.

ADVANCED DECISION TECHNOLOGY PROGRAM

CYBERNETICS TECHNOLOGY OFFICE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY
Office of Naval Research • Engineering Psychology Programs

78 10 03 034

The objective of the Advanced Decision Technology Program is to develop and transfer to users in the Department of Defense advanced management technologies for decision making.

These technologies are based upon research in the areas of decision analysis, the behavioral sciences and interactive computer graphics.

The program is sponsored by the Cybernetics Technology Office of the Defense

Advanced Research Projects Agency and technical progress is monitored by the Office of Naval Research — Engineering Psychology Programs. Participants in the program are:

Decisions and Designs, Incorporated
Harvard University
Perceptronics, Incorporated
Stanford Research Institute
Stanford University
The University of Southern California

Inquiries and comments with regard to this report should be addressed to:

Dr. Martin A. Tolcott
Director, Engineering Psychology Programs
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

or

Dr. Stephen J. Andriole
Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

2

9 TECHNICAL REPORT, TR 78-8-72

6 **RECONCILING INCOHERENT JUDGMENTS (RIJ)-
TOWARD PRINCIPLES OF PERSONAL RATIONALITY,**

14 TR-78-8-72

10 by
Rex V. Brown and Dennis V. Lindley

DDC
OCT 4 1978
F

Sponsored by
Office of Naval Research
Engineering Psychology Programs
Contract Number N00014-75-C-0426


and

Defense Advanced Research Projects Agency
Contract Number N00014-78-C-0100

12 83p.

11 July 1978

15 N00014-75-C-0426, N00014-78-C-0100



DECISIONS and DESIGNS, INC.

Suite 600, 8400 Westpark Drive
P.O. Box 907
McLean, Virginia 22101
(703) 821-2828

390 664 80 80 0- 0/4

ABSTRACT

Decision analysis involves constructing models which force logical coherence between a subject's judgments, e.g., between his choice of action and probabilities and utilities. However, it does not specify how he should reconcile any incoherent judgments. There are indefinitely many ways they can be adjusted to be coherent systems of judgment. The authors discuss two approaches for identifying one ideal set of reconciled judgments for a subject, given some or, in the limit, all potential incoherent "readings." They both call for higher order judgments bearing on the "precision" of the subject's original readings. One is a straightforward extension of Bayesian updating with the readings serving as data to update a prior. The other involves minimizing adjustments, taking into account the stability of the readings as probabilistically measured.

ACCESSION for	
NTIS	Write Section <input checked="" type="checkbox"/>
DDC	Bull Section <input type="checkbox"/>
MANUSCRIPT	<input type="checkbox"/>
DISPOSITION/AVAILABILITY CODES	
SPECIAL	
A	

78 10 03 034
ii

SUMMARY

People frequently make assessments, recommendations and decisions at all levels of national life without consciously taking into account important factors they are aware of. They rely on their first perception of a problem, when maturer consideration might suggest a very different, and generally sounder, conclusion. "Maturer consideration" may be thought of as reconciling a subject's perceptions of related issues into a coherent whole--which they may not be initially. (In surveying, the use of triangulation rather than a single pair of bearings to locate an object raises similar issues.) Informal ways of cross-checking and reconciling perceptions may work adequately in simple problems, but they often fall seriously short on complex major problems. This paper seeks systematic principles and applied techniques for reconciling incoherent judgments effectively.

In formal decision analysis, models are constructed which force logical coherence between the judgments of a given subject, S, at a given point in time. For example, his actions are made to conform to his judgments of probability and utility, through axioms that require choice to maximize expected utility. However, it is by no means clear what principles--within or beyond current decision theory--should guide this "forcing." How should an "incoherent" subject reconcile his judgments into a coherent decision model?

In general, there are many ways to adjust a subject's incoherent readings to be coherent. For example, they can all be brought into line with some minimally specified subset of readings. But how, in principle, should one reconciliation method be preferred? Does it matter whether a limited set of S's overspecified readings are being reconciled; or, in

the limit, all his potential readings? And how does S decide in advance which potential readings to take, given that it is impractical to take them all?

It seems clear that a higher order of judgment is needed from S which bears on the validity of his original readings. What form should it take?

One possibility is a prior distribution and likelihood function, which, through Bayesian updating, gives a posterior probability for some ultimate reconciled judgments. This requires no new concepts outside of current decision theory, but it is awkward to implement and runs into the problem of resolving higher order incoherence.

Another possibility is to take measures of S's cognitive stability for his primary readings--perhaps a joint probability distribution on possible shifts under further reflection. The preferred reconciled system of judgments is "fitted" to the overspecified system of incoherent readings so as to minimize some measure of stability disturbance. Although such an approach appears to map well onto intuitively reasonable informal practice, the underlying rationale is not fully developed, and specific procedures are still to be proposed.

The higher order readings might again be based on a logic quite different from decision theory--such as Zadeh's fuzzy reasoning (Zadeh 1977, Watson et al. 1978). In all these approaches incoherence in higher order judgments has to be satisfactorily accommodated.

Limited work of an exploratory and discursive nature on these issues and on possible solutions, theoretical and prac-

tical is reported here. A more sharply focused technical treatment of a special case (Bayesian updating of the probability of an event) is reported in Lindley et al. (1978).

Although one cannot say whether coherence analysis may become a major area of research within and beyond decision theory--on a par with, say, utility theory--significant activity has already been generated among researchers and university teachers in the decision theory world. (See French 1978.)

PREFACE

In this paper we discuss some general issues bearing on the reconciliation of incoherent judgments by an individual. Particular techniques have been discussed in Lindley et al (1978).

In Section 1 we introduce the problem and the main issues to be addressed, namely: is there a uniquely best reconciliation of a total psychological field; what principles should guide the reconciliation for a subset of readings on this field; what strategy should be adopted for seeking out and reconciling potential incoherence, including the strategy for choosing one or more decision-analytic models.

Section 2 discusses general considerations to be taken into account in addressing these issues and proposes a conceptual framework.

In Section 3, we explore two potential principles for reconciliation of incoherent judgments: an extension of conventional Bayesian up-dating calling for higher order assessments, and an approach based on assessments of the validity or "shiftability" of primary readings.

In Section 4, we conclude with some general observations and suggest lines for further inquiry. A glossary of terms developed for this field of inquiry is proposed.

The work described was initiated under the sponsorship of the Office of Naval Research (Engineering Psychology Programs) under contract N00014-75-C-0426 and continued under the sponsorship of the Defense Advanced Research Projects Agency (Advanced Decision Technology Program) under Contract N00014-78-C-0100 with the office of Naval Research acting as technical monitor.

The technical basis for the project was laid at University College London during 1976 while Dr. Brown was a Senior Research Fellow in the Department of Statistics and Computer Science and Dr. Lindley was Department Chairman. A large part of the material presented was developed later, during the course of extensive discussions with Professor Amos Tversky of Stanford University. His ideas have heavily influenced this report, though responsibility for any errors or misconceptions is the authors' alone. However, a joint paper addressing similar issues to those raised in this report, authored by Dr. Tversky and the present authors, is in preparation and planned for publication in a professional journal.

CONTENTS

	Page
ABSTRACT	ii
SUMMARY	iii
PREFACE	vi
FIGURES	x
1.0 INTRODUCTION	1
1.1 The Problem	1
1.1.1 Human judgment as an improvable instrument of policy	1
1.1.2 A defense intelligence application	2
1.2 Current State of the Art	4
1.2.1 Conventional approaches	4
1.2.2 Status of alternative approaches to reconciliation of incoherent judgments	5
1.3 The Role of Decision Theory	6
1.3.1 Decision theory as prescription	6
1.3.2 Decision theory as a test of coherence	9
1.4 Illustrative Cases of Incoherence	12
1.4.1 Example 1 - probability of a sporting event	12
1.4.2 Example 2 - continuous probability assessment for energy demand	13
1.4.3 Example 3 - social utility function for nuclear regulation	16
2.0 PROBLEM FORMULATION	18
2.1 Issues to be Addressed	18
2.1.1 Unique rationality	19
2.1.2 Implications of rejecting unique rationality	22
2.1.3 Partial reconciliation	24
2.1.4 The strategy for digging in the psycho-field	25

2.2	Elements of a Solution	27
2.2.1	Basic steps	27
2.2.2	The role of higher order assessments	31
2.2.3	A Bayesian updating paradigm	33
2.2.4	A Paradigm based on reading stability	35
2.2.5	Other paradigms: fuzzy reasoning, etc.	36
3.0	TOWARD AN APPROACH BASED ON "STABILITY" OF INITIAL READINGS	38
3.1	Elements of an Approach	38
3.2	Stability as a Measure of Validity	39
3.3	Evaluating Alternative Reconciliations	40
3.3.1	Comparing minimal structures before taking readings	40
3.3.2	Reconciling an overspecified system after taking readings	42
3.4	Use of Measure of Reconciliation Quality	44
4.0	CONCLUSION	46
4.1	Bayesian Updating	46
4.2	Stability-Based Adjustments	47
4.3	What Next?	47
	NOTES	49
	GLOSSARY	62
	REFERENCES	65
	DISTRIBUTION LIST	67
	DD FORM 1473	70

FIGURES

Figure		Page
1-1	Decision Theory as Prescription	7
1-2	Reconciling a Continuous Probability Assessment	15
2-1	An Overspecified Decision Tree Structure	28
2-2	Reconciliation of Incoherent Judgments - Basic Elements	29
3-1	Stability Adjusted Reconciliation	42

1.0 INTRODUCTION

1.1 The Problem

1.1.1 Human judgment as an improvable instrument of policy - Human judgment is a major resource, perhaps the only ultimate resource, at the service of national policy. It resides in administrators, scientists, and technicians, and it serves the needs of defense, government, technology, professional business practice; human judgment governs the conduct of our private lives. Poor judgments lead to poor decision making and unsatisfactory achievement of objectives.

Good judgments clearly depend on the quality of information available to the judge, and very substantial national resources are devoted to this end in the form of scientific technical research, intelligence systems, and other inquiries. Good judgment also depends, no less significantly, but much less obviously, on how the information is processed by human subjects. Much valuable research and development has been done in recent years to this end, notably in decision analysis, through such research programs as those sponsored by ARPA and ONR. The centerpiece of such approaches has been coherent structures whereby judgments of ultimate interest, for example, a choice between options, are derived logically from a sufficient set of, in a sense, more elementary judgments, say of probability and utility.

Human judgments, even for a given subject, do not typically form neat coherent structures in the sense that all potential judgments the subject might make are logically consistent with each other. A choice derived from one set of judgments, for example, may not match with a choice made from another set of judgments, even though in a perfectly coherent subject the two would coincide. However, a perfectly

coherent subject would not need such tools as decision analysis, statistical inference, or other logical tools since any conclusion he wanted to draw would automatically, by direct judgment, coincide with any other legitimate way he might analyze data available to him. Given that subjects are, in general, incoherent, often to the point of simultaneously holding highly conflicting beliefs, how can incoherent judgments be reconciled, and in the process improved?

Is there only one logical way for a subject, say a decision maker, to process what he has in his head at any point in time? Is there only one "right" decision or inference for him on any issue in terms of his total psychological field or psycho-field?

What should the demonstrably incoherent subject do if he wishes to be rational? Is there some unambiguous and compelling principle by which his incoherence can be resolved--beyond training the subject not to make obvious assessment errors due, for example, to misunderstanding the meaning of probability? (The counterpart in surveying would be to make sure the theodolite is held correctly.*)

Of course, any system of assessments can be made coherent, by arbitrarily adjusting their values. But is any one such reconciliation superior to another and on what grounds? Coherence by itself does not seem to provide an answer, and without one the very foundations of decision theory as a prescriptive tool are challenged.

1.1.2 A defense intelligence application - The effectiveness of the National Defense Effort depends critically

* See Notes

on human judgment: evaluations, predictions and assessments, choices among courses of action. What is the relative value of one weapon system compared to another? How likely is the Mideast war within the next twelve months? How many soldiers do the Soviets have under arms in Eastern Europe? Should NATO mobilize in the face of an ambiguous threat?

A U.S. defense official was faced with the problem of making his probabilistic assessment of the number of Soviet soldiers under arms in Eastern Europe as of November 1977. The "target" quantity could be expressed in a number of alternative ways, as a function of different components, for example: as the number of Soviet installations in Europe times soldiers in the Soviet army times the fraction in Europe; or as the figure for March 1976 (when an extensive study had been made) times the proportional growth since; or as the number in a particular sector of Poland (for which intelligence had a reliable estimate) divided by its fraction of the total; or in any of a number of different ways.

From survey and other sources, he made probabilistic estimates for each component in each formula and, by using statistical theory, found the implied probability distribution for the target quantity. However, the different decomposition functions yielded quite different probability distributions for the same target and, derivatively, quite different defense decisions. What distribution should he base his decision on and how could he justify that choice? If one choice is as good as another, why should he not simply make a direct assessment of the target and disregard any of the more sophisticated theoretical approaches to uncertainty assessment?

Other examples of reconciling incoherence are discussed in Section 1.4.

1.2 Current State of the Art

1.2.1 Conventional approaches - Analytical techniques have been developed--notably within the framework of personalist inference and decision theory--to help make judgments which most effectively use available information, expertise and preferences.* Primarily what they contribute is a discipline of coherence. That is, they ensure that one set of judgments is consistent with another, for example, that a preference between actions coincides with that implied by a set of uncertainty and value judgments, measured by probability and utility; or that a directly assessed probability is consistent with indirect assessments combined by probability theory (say in the form of Bayes' Theorem). However, if the judgments derived in different ways do not coincide, i.e. the subject is incoherent, the theory does not say how to resolve the incoherence. It is because subjects are incoherent that they need analytic aids in the first place.

When different analytic approaches (for example, different decision or inference models) yield different answers (that is, implied judgments), even though they are based on the same body of expertise and information, there is a serious practical problem, not only of determining which conclusion to draw, but in justifying to third parties the validity of this conclusion. Techniques for performing a sound reconciliation are needed as is a theoretical base for justifying those techniques. Neither is supplied by current decision theory.

What is called for is a technology, firmly grounded in normative and descriptive theory to improve the

* See Notes

quality of judgment and decision making, by making the most efficient use of the totality of data, however conflicting, available to decision makers.

1.2.2 Status of alternative approaches to reconciliation of incoherent judgments - Modest steps have been taken to initiate research on new approaches to this problem, primarily philosophical and mathematical rather than behavioral. (See Brown and Lindley 1977; and Lindley et al. 1978).

The first area of research examines the fundamental philosophical principles according to which practical techniques of judgment reconciliation should conform. They are still by no means clearly established. One view is that higher order models using no more than the standard Savage axioms of decision theory will suffice (Savage 1972), for example, treating raw incoherent judgments as data to be processed "Bayesianly" by eliciting special priors and likelihood functions. Another is based on the concept of assessment validity for each element in a structure of incoherent assessments from which a "most valid" reconciled system of assessments can be derived. A third view is that the concept of fuzzy sets can be used; and there may be further formulations worth exploring.

The second area of research involves developing specific mathematical formulations for special cases of assessing event probabilities within the general "Bayesian updating" paradigm. A general model for the analysis of probability assessments is introduced, and two approaches, called internal and external, to the reconciliation problem are developed. In the internal approach, one estimates the subject's "true" probabilities on the basis of his assessments. In the external approach, an external observer updates his own coherent probabilities in the light of the assessments made by the subject. The two approaches are illustrated

and discussed. Least-squares procedures for reconciliation are developed within the internal approach.

So far our discussion has been largely in terms of probabilities. Procedures have been developed by Novick and Lindley (1978) for resolving inconsistencies in utility functions for a subject which have been derived in different ways. The two curves have a third curve fitted to them by least-squares, and this is presented to the subject for his evaluation. He subjects it to plausibility checks (that is, he is implicitly calling up more material from his external field to check its plausibility and then to modify it accordingly).

1.3 The Role of Decision Theory

1.3.1 Decision theory as prescription - Decision theory as a prescriptive paradigm for decision making by a subject has been grounded by Savage and others in axiomatic conceptual systems whose essence is to derive the logical action implications for a subject of certain persuasive behavioral axioms (Savage 1972).

To put it at its simplest, Figure 1-1 shows a prototypical prescriptive use of decision theory. The target judgment, that is, the one to be prescribed, is whether action A is preferred to \bar{A} . This judgment can be inferred from another set of judgments, for example, from probabilities and utilities organized on a decision tree, as shown. The primary readings, in this case p , u_1 , u_2 , u_3 , on the subject (call him S) are then the probabilities and utilities, and the derived judgment is a binary variable--whether A is or is not preferred to \bar{A} . If the q readings are minimally specified, as they are in this example, there is no possibility for incoherence, and the most obvious conclusion from the




Figure 1-1
DECISION THEORY AS PRESCRIPTION

exercise is that we should prescriptively prefer \bar{A} to A if that is what the readings imply. However, as soon as we take a direct reading on whether A is preferred to \bar{A} , we have an overspecified system of reading and potential for incoherence.

A naive, but by no means rare way for handling this possible embarrassment is to take only a minimally specified set of readings, for example, take no more readings than are necessary to complete an appropriate decision tree. In that case, the issue of incoherence never surfaces. However, this bars one from the approach, which appeals greatly to common sense, of addressing issues of judgment in several different ways, which might be incoherent.

A more common position taken, at least implicitly, by some decision analysts is that some readings take precedence over others which are either disregarded or forced into line if there is incoherence. Thus, it might be argued that the probabilities and utilities on the left of Figure 1-1 are somehow more valid than the direct choice on the right and, therefore, the subject "should" make the derived choice. A common variant of this position is that readings in complex formulations take precedence over more simple ones. For example, a choice based on a decision tree is preferred over a choice made directly.

If a directly assessed posterior distribution differs from one derived by updating a prior with a likelihood function, the practice would be to accept the latter and disregard the former. At an informal level this practice could clearly be questioned if, for example, one had serious doubts about a subject's ability to assess a "reliable" likelihood function or an uncontaminated prior (see Brown, 1969).

This is not to say that experienced decision analysts would subscribe to this principle if it were called out explicitly to them, but it would seem to characterize much decision analysis that is actually done, and to be honest it is consistent with how we used to teach it at business schools in the early '60s. Of course the problem would not arise if subjects were coherent, but then who would need decision analysts!

There is nothing objectionable about the notion of precedence between readings, as we shall see later. It is clear that some judgments deserve to be taken more seriously than others, say because the subject has more familiarity with that type of assessment. What is not obvious is that the less valid readings should be disregarded entirely or that the readings called for in more complex models automatically take precedence.

It used to be argued that a posterior derived from a prior and from a likelihood function is to be preferred because "people can't do probability theory in their heads." We have argued elsewhere (Brown 1968), as have de Finetti and Good, that if the likelihood function calls for unfamiliar hypothetical assessments and the prior is contaminated by knowledge of the updating data, the derived posterior may be less "valid" than one directly assessed (which might be soundly based on informal extrapolation from past posterior assessments in comparable situations and validated by hindsight). A fascinating example of this phenomenon in astrophysics has been reported by Sturrock (1973).

1.3.2 Decision theory as a test of coherence - An alternative conception of the role of decision theory is that it tests overspecified readings (on probability, utility and choice) for coherence, where coherence is derived from the

usual (e.g. Savage) axioms of decision theory. It is generally held that coherence is a "good thing" leading to valid expectations of higher utility to the subject than behavior not so blessed. Indeed one of the authors of this paper (Lindley 1973) has been known in the past to declare that "coherence is all!". It certainly seems intuitively plausible as a loosely expressed proposition, though we do not know of its being given an unambiguous interpretation or being compellingly demonstrated, and we will not attempt to do so here. However, it is by no means clear how coherence is to be assured within the tenets of decision theory. It may well be that after all "coherence is not enough!".

This problem was recognized by Savage himself (Savage 1972):

Logic, to which the theory of personal probability can be closely paralleled, is ... incomplete. Thus, if my beliefs are inconsistent with each other, logic insists that I amend them, without telling me how to do so. This is not a derogatory criticism of logic but simply a part of the truism that logic alone is not a complete guide to life. Since the theory of personal probability is more complete than logic in some respects, it may be somewhat disappointing to find that it represents no improvement in the particular direction now in question.

A second difficulty, perhaps closely associated with the first one, stems from the vagueness associated with judgments of the magnitude of personal probability. The postulates of personal probability imply that I can determine, to any degree of accuracy whatsoever, the probability (for me) that the next president will be a Democrat. Now, it is manifest that I cannot really determine that number with great accuracy, but only roughly. Since, as is widely recognized, all the interesting and useful theories of modern science, for example, geometry, relativity, quantum mechanics, Mendelism, and the theory of perfect competition are inexact; it may not at first sight seem disquieting that the theory of personal probability should also be somewhat inexact. As will immediately be explained, however, the theory of

personal probability cannot safely be compared with ordinary scientific theories in this respect.

I am not familiar with any serious analysis of the notion that a theory is only slightly inexact or is almost true, though philosophers of science have perhaps presented some. Even if valid analyses of the notion have been made, or are made in the future, for the ordinary theories of science, it is not to be expected that those analyses will be immediately applicable to the theory of personal probability, normatively interpreted; because that theory is a code of consistency for the person applying it, not a system of predictions about the world around him.

Is there some enriching of the Savage axioms which can uniquely force coherence; or can the axioms do it after all, Savage's own disclaimers notwithstanding? We have so far no satisfying answer to this question, though we address it obliquely later in this paper.

Short of resolving this fundamental philosophical issue, one is led to ask what defensible principle can be used to guide the forcing of coherence?

The resolution of incoherence is a familiar dilemma among decision analysts. A common principle has been to put the responsibility on the subject to produce reconciled assessments. The role of the decision analysis then is to draw the subject's attention to the fact that, for example, probabilities do not add up to unity and have him go away and come back when he has made them add up to one. There is nothing in principle objectionable to this procedure. A primary function of decision analysis, after all, is to replace a single complex indigestible problem by a logically equivalent set of more manageable problems. However, it does not address the question of how, if at all, a modified set of judgments (for example one based on forcing coherence) is better. Nor could it, of course, without establishing what "better" means.

1.4 Illustrative Cases of Incoherence

To give concreteness, let us consider some examples where the logic of decision theory (including probability and utility theory) might detect incoherence between over-specified readings. They involve target judgments with alternative ways of assessing them directly or indirectly. Each way calls for a minimally specified set of readings, and between them they represent an overspecified set of readings with potential incoherence.

1.4.1 Example 1 - probability of a sporting event -

Suppose the subject's target judgment is the probability of Cambridge winning the next Boat Race. There are several ways we could start to help the subject make his target assessment by digging different numbers out of his cognitive field. We could have him make a direct assessment of that probability. He could then indirectly assess the target by using Bayes' Theorem on yesterday's (prior) probability and today's information that the regular Oxford cox is sick. Or he might assess the target with a conditioned assessment model where rain is the condition event, that is,

$$P(\text{Cambridge}) \equiv P(\text{Cambridge} | \text{rain}) P(\text{rain}) + P(\text{Cambridge} | \text{no rain}) (1 - P(\text{rain})).$$

What one has then is a set of models which share a characteristic that they each imply the value of a target variable such as the probability of Cambridge winning. In general, one expects some of the models as quantified to yield different target values from others and, therefore, to have demonstrable incoherence with each other.

This special case of reconciling probability assessments for an event is covered in some detail in Lindley et al. (1978).

1.4.2 Example 2 - continuous probability assessment for energy demand - The assessment of probability distributions for a continuous variable represents a more complex reconciliation task. A real life example of this kind in which one of the authors (Brown) was involved provided a major motivational stimulus to develop the reconciliation of incoherent judgment as an area for research and technical development.

A senior member of the staff of the Federal Energy Administration was charged with presenting to Congress probabilistic estimates of energy demand, broken down by area, form and end use. Substantial survey and other empirical and analytic work had been done in the field, and DDI was asked to help produce defensible assessments from the available evidence.

Extensive experience in the survey field suggested that the best way to attack this, like most other estimation problems was to attack them from a number of directions and "pool" the results. (See Brown 1963, pages 375, 376).

In this case, we had available a large number of different ways of making any particular estimate. For example, the demand for lighting energy in schools in the Northwest could be expressed as a number of different "target functions," that is, expressions which give some target quantity as a function of two or more arguments.

The demand could be expressed in terms of the number of students, bulbs per student, hours per bulb and average wattage per hour. The arguments in such a function could be estimated from available surveys, censuses, published statistics, and engineering studies. Alternatively, it could be expressed as the demand in 1968 (for which a substantial SRI study was available), times the change since

1968 (which could be based on an informal evaluation of economic and demographic trends). Yet again it could be based on an intensive survey of lighting per student in the Washington, D.C. area in conjunction with a judgmental assessment of how Washington and the Northwest differ, and statistics on the number of students in the Northwest.

Up to a dozen approaches of this kind were available for any particular assessment and, taken one at a time, they produced very different numbers, sometimes differing by a factor of three or more, even when the conflicting estimating approaches used appeared individually reliable.

By assigning probabilities to each of the arguments in each of the target functions (more generally a joint probability distribution), a probability distribution for the target variable could be routinely derived (see Brown 1969, Chapter 9). However, as one might expect, the target distributions did not coincide--in fact were widely different--indicating incoherence in the input distributions. Figure 1-2 gives a simple illustration with just two decompositions giving derived assessment L' and L'' .

To present these conflicting assessments to the public could be politically embarrassing to the FEA to say the least. Past practice had often been to present whichever assessment was thought to be "best" and to suppress the others. An alternative, which we had on occasion used in the past in similar situations was to mechanically "pool" the several independent estimates in a way which reflected the relative dispersion of the derived assessments. For example, the pooled mean could be a weighted average of the component means, with the weights inversely proportional to the variances, the pooled precision (reciprocal of variance) being estimated as the sum of the component precisions.




Figure 1-2
RECONCILING A CONTINUOUS PROBABILITY ASSESSMENT

This procedure had some intuitively desirable properties such as locating the final mean closer to those component estimates in which there was most initial "confidence" as measured by variance. On the other hand, the pooled precision formula took no account of how far apart the component distributions were nor even whether they overlapped, which they often did not. Moreover, it would clearly be more satisfactory--for example, when presenting testimony to Congress--to be able to present a coherent system of probabilistic assessments giving a single reconciled target assessment \hat{L} , if one but knew how.

Moreover, it was politically important to have a procedure for assuring coherence that could stand up to academic scrutiny. Informal adjustments of the component assessments would not meet this criterion, and we came to the realization that we knew of no generally accepted procedures which would. This provided us with the practical motivation to seek a reconciliation procedure with a solid theoretical basis.

1.4.3 Example 3 - social utility function for nuclear regulation - A government agency was in the process of determining what depositories for nuclear waste would be acceptable for purposes of regulation. The agency arranged to have constructed a utility function intended to represent the nation's values in terms of costs and different types of radiological health hazards. The national utility function was to be derived from individual utility functions elicited from selected individuals, including a senior official of the agency.

He was asked, for example, to assess his personal trade-offs for: equating numbers of instant deaths; lingering cancer fatalities; female sterilities; and sub-normal offspring

in the U. S. population, now and in generations yet unborn. He found that different pairings of these grave potential consequences produced seriously inconsistent trade-offs. For example, he indicated indifference between: one death and twenty cases of sterility; between one sub-normal child and two deaths; and between one sub-normal child and 1.1 sterilities. Note that by combining the first two equalities, we would deduce that he considered one sub-normal child was the equivalent of forty sterilities rather than the 1.1 which he evaluated directly, clearly a major inconsistency.

In the process of informally reconciling this incoherence, it appeared that the official had been focusing on two quite different aspects of sterility with the two pairs of judgments that involved it. When equating one death with twenty sterilities, he was considering primarily the broad sociological implications of reduced fertility such as the control of population growth, which made it appear not too serious. When equating one sub-normal child to 1.1 sterilities, he was thinking primarily of the personal anguish that sterility might cause an individual. When the incoherence and its sources were brought to his attention, he made sure that both aspects figured comparably in all pairings, and informal reconciliation was fairly painlessly achieved.

2.0 PROBLEM FORMULATION

2.1 Issues To Be Addressed

There are essentially three general issues which concern us in the area of reconciling incoherent judgments (RIJ):

1. Is there any such thing as unique rationality, in the sense that there is one best way to reconcile all potential readings in a subject's psycho-field?
2. How should a subject reconcile any given, partial set of incoherent readings?
3. What strategy should he adopt in taking readings, i.e. for "digging in the psycho-field"?

Each of these items has three facets: normative, psychological, and applied. The normative question concerns the manner in which incoherence should be resolved. The psychological descriptive question concerns the manner in which people actually resolve incoherence. The applied problem deals with the implementation of procedures for the resolution of incoherence. Clearly, the applied problem is closely related to both the normative and the descriptive problems.

We hope to propose and develop one or more solutions to these problems which are theoretically sound and which have promise for practical implementation.

In all cases we are concerned with the psycho-field of a single subject to whom all elicitations and assessments

* See Notes

refer, at a single point in time and predicated on a fixed body of information, the history of sensory data received.

A satisfactory theory of reconciliation is likely to have implications to several related problems, such as: the problem of amalgamating experts' opinions (Morris 1974), the problem of forming a subjective probability function for a group of individuals, and the problem of defining the value of a decision analysis (Watson and Brown 1978).

Note that issue three, the strategy of digging in the psycho-field, may have a place even if no partial incoherence has been found initially. It is motivated by the expectation of incoherence after digging.

2.1.1 Unique rationality - The idealized rational, coherent subject is prepared with a probability, a utility and a choice for every conceivable circumstance; and all these values cohere in a unified system which obeys the rules of the decision theoretic calculus. Real man, or at least the subject as measured by available probability and other elicitation instruments, is incoherent. Is there inside real man a coherent man that he would wish to be--and is this coherent one unique? In particular, for any event, A, is there a unique coherent, rational probability, $P(A)$ --and if so, how can it be determined? If there are such unique values, then what properties, for real man, do they have other than coherence, or is coherence, after all, all?

Even if there is no unambiguous principle according to which initially incoherent readings can be reconciled, can we at least identify some criterion according to which one reconciled system of assessment is preferred to another?

At a weaker level still, is there any priority between possible reconciliations? If not, it would appear there are no defensible grounds for favoring any choice, probability or judgment that might be attributed to a subject over any other.

Clearly there is some priority possible in reconciled assessments system space (that is, possible true values) since there are certain regions that no readings suggest. For example, if the different ways of getting the probability of Cambridge winning the Boat Race all lie within the range .3 to .5, we can throw out for further consideration any values outside that range.*

If the ultimate reconciliation is no more than constrained to the "obvious" region, this has some alarming implications. It would appear to remove any motivation for improved rationality since any way of getting at a target judgment would be as good as any other, and the decision theorist would cease to have any prescriptive role.

There is a good deal of intuitive appeal to the notion that there is one "right" way to process the totality of a subject's information, judgment and perception of a subject at a particular point in time, which produces a single "best" target assessment. It is appealing to think that if only we applied infinite and impeccable pains to the analysis of that corpus of knowledge, one would get the "right" answer.

An alternative interpretation of unique reconciliation is a model of a perfectly rational subject who starts as a

* See Notes

fetus with some basic judgmental set (including priors and likelihood functions), and who uniquely updates it through life in a way determined by the sensory data he receives. One might allow this updating also to be influenced by changing human chemistry, which could autonomously change his tastes and therefore his utility judgments. Of course, part of the chemistry is imperfection in neural connection which leads to irrationality.

However, if one considers only probability judgments for a moment, it would be tempting to imagine that one emerges from the womb with some kind of uniform prior joint distribution over everything the world has to offer including new sensory data. One updates this prior by using Bayes' Theorem as data impinges on one's senses and their probabilities increase to one.

One probably need not argue for any particular interpretation of the uniform prior, nor even how any fetal incoherence were resolved, since the accumulation of lifelong experience would soon make posteriors very insensitive to the choice of initial prior. This logic would appear to cover all eventualities, including the updating of likelihood functions which are part of the subject's dynamically updated joint probability distribution.

Take several rational subjects hearing a radio announcement that life has been found on Mars. The infant will have a uniform likelihood function (that is, undiagnostic) since the message is incomprehensible until he learns English. The child who has not yet learned to be skeptical has a highly peaked likelihood function. And so on.

One might argue on grounds of intuitive plausibility that, to a decent approximation at least, a subject's topical impeccable assessments are a close-to-determinate function of all that the sensory data he has ever perceived. If any remaining ambiguity is accounted for by the chemistry of the subject, then there is only one set of conclusions a subject can rationally hold at any one point in time, and the notion of unique rationality would appear to be sustained.

2.1.2 Implications of rejecting unique rationality - If we do not accept that there is a unique rational assessment system for a given subject at a given time, what are the implications? Do we lose any justification for attempting some rationality if there is no ideal towards which one can, conceptually at least, aspire? If there is not a single rational assessment, might there not be a set, perhaps a fuzzy set, of rational assessments, all of them equally acceptable, which excludes at least some of the systems one might have started with?

The weakest case of this would be the set of all coherent systems. Any single element in the system would be free to take on any value, but there would be a limited number of degrees of freedom which would impose some constraint on other elements.

The idea of there being a single correct analysis for a subject is critical to any evaluation of a proposed, necessarily imperfect analysis (see Watson and Brown 1978). The direct value of analysis (as opposed to indirect values such as improved communications or psychological peace) depends on the fact that, left to himself, a potentially incoherent subject will come up with a choice or an assess-

ment which differs from the "correct" one. The difference in expected utility (according to correct probabilities and utilities) between the action he will choose and the action he would have chosen given perfect rationality can be interpreted as "the cost of irrationality" (see Brown et al. 1974, p. 359). The expected cost of irrationality, thus defined, will give a measure of the value of perfect rationality (analogous to the value of perfect information*).

If the existence of a perfect analysis for a given subject (specifying the person, the time and the information received) is denied, a value might still be imputed to a proposed piece of analysis. However, the task of conceptualising it is certainly much greater, if there is no benchmark or anchor point corresponding to perfect analysis to scale the value.

If one takes the position that there is no sense in which one action implied by one internal model has precedence over any other, actual or potential, then clearly no analysis has any value--an intuitively quite unacceptable conclusion to those of us who make their living doing decision analysis!

It is possible, however, that some position weaker than the assertion of perfect analysis is possible, as we have suggested. If one posits that there is a set of plausible candidates for the role of perfect analysis each with a different measure of strength attached to it, then one could take a weighted average of perfect analysis values predicated on each of them being the perfect analysis weighted appropriately. This, however, smacks of the dreaded blight of "ad hockery"! It may, however, stand as a suitable Aunt Sally until knocked down by some intellectually more satisfying approach.

* See Notes

2.1.3 Partial reconciliation - Let us characterize all potential readings, in principle incoherent, of a subject, S, as Q, whose unique reconciliation is π . Let q be any subset of Q, for example, readings that have actually been taken (say for one or more minimally specified models). $\hat{\pi}$ is an estimate of π , itself coherent, based on q. The process of going from q to $\hat{\pi}$ we might call partial reconciliation.*

The major practical task of RIJ (reconciling incoherent judgment) is to find an implementable procedure for partial reconciliation. However, there is still a purely conceptual problem of defining what would constitute an ideal partial reconciliation for a given q. Is it exactly the same problem as specifying a unique reconciliation for all potential readings Q? In that case ultimate reconciliation would simply be the limiting case of partial reconciliation.

A fairly obvious (but cumbersome) approach to a Bayesian would be to require a higher order, already coherent probability distribution over q and π (implying for example, a prior over π and a likelihood function for π given q). Interfacing actual readings q with these higher order probabilities immediately gives a conditional distribution of π given q. A conditional expectation of π given q would then give us a partial reconciliation $\hat{\pi}$ as required.

A mechanistic approach which does not conform to intelligent informal practice would be to pool target estimates, that is, weights to points in "target space" which are non-zero whenever there is at least one way of modelling subject's assessments to produce that value. It is not clear how one

* See Notes

assigns variable non-zero weights other than by some measure of "validity." (See Section 3.0.) The fall-back position of course would be equal weights and one might then simply take an unweighted average (if the target is a scalar like the probability of Cambridge winning); or the center of gravity (if the target is a vector, e.g., a probability distribution); or a least-squares fit.

2.1.4 The strategy for digging in the psycho-field - Faced with making a target judgment, the subject first has to decide what readings to take. He can take a direct reading on the target judgment; that is he can ask himself directly which act he prefers or what his target probability is. He can make the target judgment indirectly by taking readings on a minimally specified structure such as a decision tree. Or he can take readings on an overspecified assessment structure, such as two alternative decision trees for the same choice. What should he do?*

Intuition and analogy with triangulation in surveying suggests strongly that the quality (whatever that may mean) of the target judgment will be enhanced as one takes several "bearings" on the target; that is, as the subject extends the conversation to include more and more of his psycho-field, notably by taking readings on more and more overspecified assessment structures. Our informal practice is certainly to look at a knotty problem in a number of different ways in the hope of converging on some kind of "solid" conclusion. In the limit, if we had the time and patience, we would consider everything we could think of that had a bearing on the problem at hand; and if we knew how to do it right we would presumably have the ultimate reconciliation we have sought earlier.

* See Notes

Presumably some measure of judgmental quality is expected to improve as we dig further and further into the psycho-field. In a practical situation, quality has to be traded off against the increased cost and delay of so doing.

What would be an appropriate measure of judgmental quality and how can we predict it as a function of alternative strategies of digging? If we have the tools to achieve this, we are left with a conceptually straightforward optimization task.

Improvement in the quality of judgment due to reconciliation is presumably related to the amount of incoherence to be reconciled, thus some measure of it is needed.*

One approach to attacking the issue would be to examine common informal practice among intelligent subjects and probe to see whether there is some defensible rationale behind what they do. A subject would be asked to assess the probability (or other target) in question; he would then be asked why he made that assessment. Commonly, one or more of the standard indirect probability models (such as conditioned assessment or decomposed assessment or Bayesian updating) will be present in S's more or less conscious awareness. Implicitly he is using these models, and the exercise is largely to have him do so explicitly and confront the two or more findings.

If the findings differ, the subject is given information on what changes in his input judgments would reconcile the models, that is he needs to raise this prob-

* See Notes

ability or lower than one. Thus, explicitly or implicitly, a space of acceptable reconciling adjustments is defined. The subject takes his pick. Now the key question is: why does or should the subject take one set of adjustments rather than another? That is the major issue of our formal enquiry.*

2.2 Elements of a Solution

2.2.1 Basic steps - The process of taking (partial) readings and reconciling them partially ultimately involves some basic steps which can be illustrated in the context of the decision tree example given in Figure 1-1.

In Figure 2-1 we take the minimally specified tree of Figure 1-1 and make it overspecified. In other words, we add assessments which could be inferred from assessments already made if S were coherent. The assessments are now potentially incoherent. At least four minimally specified systems of readings can be constructed from those marked, each of which could imply a different target judgment on whether A is preferred to \bar{A} , as shown in Figure 2-2.

In the context of a particular target judgment T (in this case whether act A is preferred to \bar{A}), one or more target functions* are specified, each of which gives a derived reading q'. In this case there are three target functions:

- o the choice could be assessed directly (the target function as an identity);

* See Notes




Figure 2-1
AN OVERSPECIFIED DECISION TREE STRUCTURE




Figure 2-2
RECONCILIATION OF INCOHERENT JUDGMENTS—
BASIC ELEMENTS

- o it could be calculated from a comparison of two utility readings; or
- o it could be based on a fuller decision tree with an event probability and conditional utilities.

Any set of possible elicitations is an assessment structure, and the arguments in a target function represent a minimal assessment structure.* (It is, however, possible to have assessments which do not appear in any target function*).

Any assessment can have at least four types of value. It can be a primary reading q ; it can be derived from one or more target functions based on minimal readings (q' , q'' etc.). It can be a perfect assessment based on ultimate reconciliation of the subject's total psycho-field (π). Or it can have one or more partially reconciled values based on one or more sets of overspecified readings (π , $\hat{\pi}$, etc.).

Note that the overspecified assessment structure may include assessments additional to those required by target functions. They could involve probabilistic relationships between arguments within and across target functions. These could be additional sources of potential incoherence and would need to be taken into account in the process of reconciliation.

The issues raised in Section 2.1 above reduced to defining π (ultimate reconciliation), deriving π from q (partial reconciliation) and selecting an assessment structure for q ("digging" strategy).

* See Notes

2.2.2 The role of higher order assessments - There appears to be no way to achieve either ultimate or partial reconciliation without making assessments over and above the readings to be reconciled. The formal interpretation of these higher order assessments is at the heart of the philosophical and practical problems we address.

Since the emerging reconciliation is to characterize the subject, it is clearly not appropriate to use any assessments idiosyncratic to any outside observer. In this respect the situation is different from one in which a subject is updating his belief in the light of someone else's opinion (French 1978). However, it may be convenient to treat the required higher order readings on the subject as partitioned off from the primary readings q and modeled as a separate investigator, N .

If N is modeled as a coherent probability assessor (N for normative), we have reduced the problem to one for which there is at least one closed solution. If N can produce any required probability assessments and have them cohere with each other, he can assess a prior on π and a likelihood function of π given q and derive a posterior on π .

However, this solution raises two serious theoretical and possibly practical issues. The first is: how do we address the fact that the higher order assessments will not, at least at first reading, be generally coherent? A model that supposes N to be coherent will therefore be inaccurate in a way that strikes at the heart of our problem.

Two possible approaches suggest themselves. One is to argue that the second order readings can be reconciled,

in principle, by higher order readings in an infinite regress. The first order reconciliation, the one we care about ultimately, is progressively less sensitive to how higher orders of reconciliation are performed such that the nature of the highest order reconciliation can be disregarded. Whether one can argue that such convergence holds, either invariably or under special conditions, requires analytic and psychological enquiry.

..... The second approach would be to allow the second order reconciliation to be arbitrary and treat the first order reconciliation as being therefore non-unique. Each possible second order assessment system generates a different first order reconciliation (partial or ultimate). The feasible order of second order reconciliations therefore induces a new feasible region of first order reconciliations which hopefully is more restrictive than initially. That is, the second order assessments have achieved some measure of first order reconciliation. The second approach would appear to reduce to the first if we continue the process with successively higher orders of assessment. It still remains to be established whether the first order reconciliation thus induced converges to a single point in reconciled system space.

The other bothersome issue here is the arguable assumption that there is a unique. Only if there is can we comfortably talk of priors, likelihood functions and posteriors and, more generally, joint probability distributions involving π and q . It is not quite clear if we can define π as the limiting case of partial reconciliation as primary readings are indefinitely increased, without logically unacceptable circularity.*

* See Notes

2.2.3 A Bayesian updating paradigm - One approach to the resolution of incoherence involves positing an investigator, N (distinct from subject, S), who is to determine a unique, rational system of assessments for S. Unlike S, N is treated as perfectly coherent. He has a prior distribution on S's ultimate target assessment, T, and a likelihood function for T, given S's raw readings q. Through Bayes' Theorem he can derive a posterior distribution on T.

It is probably mathematically demonstrable that the variance of N's posterior on T gets smaller, possibly to the point of vanishing, as q is extended to include more and more of the subject's potential readings Q (how fast depends on the diagnosticity of the likelihood function). This much is investigator-independent and confirms one's intuitive conviction that it pays to address a target assessment in as many different ways as possible (much as it pays a surveyor to take many different bearings on a location).

However, we are left with the problem of having a reconciliation procedure dependent on characteristics attributed to the investigator. Where do N's priors and likelihood functions come from? Since N is a hypothetical construct, they should not be idiosyncratic to N but should somehow be descriptive of S.

In this respect our problem contrasts sharply with that of updating one's belief in the light of someone else's opinion, a topic that has many formal similarities and has also been addressed through Bayesian updating (Morris 1974, French 1978).

Can we treat N as a partition of S as a coherent assessor for this purpose? Can any incoherence here causing second order fuzziness perhaps be disregarded? If the

likelihood function is informative enough, any "uninformative" prior, however defined (and therefore however reconciled from initial incoherence) may lead to virtually indistinguishable results and so be acceptable. But according to what principle should N (or S) construct the likelihood function?

This Bayesian updating approach, whereby the initially incoherent readings are treated as data which update the subject's super ego N's prior with the help of N's likelihood function, is the easiest one for a regular decision theorist, especially a Bayesian theorist, to visualize.

A special case of this approach has been developed in Lindley et al. 1978. In particular it considers the reconciliation of event probability assessment. It shows, for example that in a simple, but not implausible case, the precision of a target judgment (reciprocal of variance of posterior π) increased by a factor of three when a single minimal assessment structure (direct assessment) was added to a different assessment, that is, made overspecified by introducing a target function.

An advantage of this general Bayesian updating paradigm is that it invokes no new theory outside of the regular axioms of decision theory. However, it is not clear whether it can resolve the problem of secondary incoherence at the prior and likelihood levels or the problem of defining the ultimate reconciliation π . Furthermore, it is not clear that from a practical point of view it leads to operational procedures for partial reconciliation. The elicitations required appear to be awkward in the extreme and may not even be obtainable in principle, but the questions to which answers are needed are not unreasonable.

2.2.4 A paradigm based on reading stability - An alternate paradigm for reconciling incoherence involves modeling what intelligent people seem to do when they reconcile incoherence, rather than extending an established formal procedure such as Bayesian updating. Descriptively what happens when an averagely intelligent subject attempts to reconcile incoherent judgments?

Let us say that his target judgment is the probability distribution of lighting energy demand, a real case (referred to in Section 1.4.2) which had a large role in motivating our investigation. One of the two target functions for energy demand shown in Figure 1-2 is based on extrapolating a past estimate to the present. Let us say his expected value for 1968 demand was 1 billion kwh, and his expected value for the increase since then is a factor of two. His expectation for 1978 demand is then 2 billion kwh (only approximately if there is dependence), and his distribution about that expectation is calculated from his joint distribution on the two arguments. And let us say it produces a 95% credible interval of plus or minus 20%. He now overspecifies his assessment structure by adding the second target function in Figure 1-2 based on number of users. Let us say his expectation of the number of users is one million, of bulbs per user is two, of hours per bulb is a thousand, and of average kilowatts is fifty. His expectation of the product will be approximately 4 billion kilowatt hours, and let us say the 95% credible interval works out to be plus or minus 40%. Notice that the two derived distributions for the two target functions barely overlap, so there is substantial incoherence.

When the subject has this incoherence drawn to his attention, he might do two things (after checking for any

obvious error in individual elicitations). He could first consider how the six sets of readings (probability distributions) could be adjusted so that they cohere; that is, one or both of the distributions for target function one could be shifted up, and/or one or more of the four distributions for target function two could be shifted down. Secondly, he might see which of the readings has most "give" and in which direction. Then he "jiggles" or adjusts the readings in a way that as Dawid* has suggested minimizes "tension." And by and large the greater the incoherence he has had to reconcile, the less "firm" he feels his partial reconciliation to be and the more inclined he is to seek more potential incoherence to be reconciled by adding new target functions.

If this informal procedure makes sense, one might seek a more formal procedure which he adopts implicitly and which can be turned into a prescriptive principle, based somehow on the validity or stability of the primary readings. How if at all would such a principle relate to the Bayesian updating paradigm discussed in Section 2.2.2 above? Are they logically equivalent at some level?

In Section 3.0 below we discuss how such a codification might proceed and what kind of logical basis it might depend on. The discussion should be considered as very tentative at this stage.

2.2.5 Other paradigms: fuzzy reasoning, etc. - It is possible that the burgeoning field of fuzzy and approximate reasoning developed by Zadeh (Zadeh 1977) and others may be adapted to the reconciliation problem. There are significant current efforts to adapt it to decision analysis (Watson et

* Private communication, December 1976

al. 1978). It would seem worth exploring in the context of reconciliation.

This approach still involves higher order assessments (say characterizing readings according to a membership function of a fuzzy set), but it may be a quite different assessment from the other we have discussed.

There are other approaches for identifying a point in reconciliation space which do not involve higher order assessments, for example, a least-squares or other mechanical fitting procedure. This may prove the most immediately useful approach by reason of its simplicity of application, but it would appear to clearly disregard information that a subject would want to take into account and does when intelligently handling the problem informally.

3.0 TOWARD AN APPROACH BASED ON "STABILITY" OF INITIAL READINGS

3.1 Elements of an Approach

It is intuitively appealing to argue that some of the subject's "raw" readings are, in some sense, more "valid" than others and that this relative validity should somehow be taken into account when reconciling initial assessments.*

At the very least, it appears compelling to suppose that there is some way to assign priority between different direct readings. But how do we characterize this priority? An apparently relevant notion here is that of "firmness" with which a judgment is held. Thus, we are all firmer about $p(A) = 1/2$ where A is the event "heads" on the toss of a coin than where A is the event of Cambridge winning the Boat Race.

A satisfactory procedure along these lines would appear to have two elements:

1. a definition of the validity of initial readings;
2. a way of devising a "quality" measure for alternative reconciliations based on reading validity.

The preferred reconciliation or reconciliation method for a particular set of readings would then be a fairly straightforward optimization problem.

* See Notes.

3.2 Stability as a Measure of Validity

A suggestive approach to defining assessment validity would be to define it as a measure of analytic stability. The .5 probability of heads on a coin toss rates higher than the .5 probability that Oxford will win the Boat Race. One does not expect further reflection to shift the probability in the former case.

Initial readings could be characterized by a probability distribution on "shift on further reflection" (not with further information, which is quite another issue*). More generally, this would be a joint probability distribution reflecting for example, "shift" dependence between readings.*

The more "valid" an assessment, the more peaked its stability distribution. Some measure of dispersion such as variance of coefficient of variation would give a measure of validity--for example, for weighting purposes--but no logical priority is apparent.

An irksome problem here is how to define "shift on reflection": how much reflection and of what kind (infinite? impeccable?); how to avoid taking for granted the optimal resolution of incoherence which the "shift on reflection" itself is to be an instrument in discovering. Even if there is some degree of circularity in definition, perhaps the "further reflection" can be specified as an uncertain procedure whose expected impact can nevertheless be assessed.

Some allowance must also be made for the subject, S, being incoherent in assessing his stability distribution. Possibly any incoherence in second order elicitation of S's incoherent view on reading stability could itself be taken into account by third order Bayesian updating to make a hybrid Bayesian stability approach.

* See Notes

3.3 Evaluating Alternative Reconciliations

The primary difficulty in this approach is to find a defensible way for selecting one from among all possible reconciliations, that is, picking a point in reconciliation space.

For example, the reconciliation space might correspond to all possible values for $\hat{\pi}$ in Figure 3-1. If a quality measure could be assigned to each point in reconciliation space, then we would simply have to optimize over that space. How do we obtain such quality measures?

3.3.1 Comparing minimal structures before taking readings - With a single minimally specified assessment structure there is a straightforward first step. The target (or targets) can be expressed as a function of all the elements (since they are minimally specified, there will be only one function). The stability of the function can be calculated from the joint stability of the arguments by using the theory of the distribution of a function of random variables. The precision of this derived distribution would then appear to be a promising measure of the target function and of the assessment system that provides its arguments.

The task of choosing among alternative target functions based on minimally specified readings would then be solved. If, for example, you wanted to choose between assessing a posterior directly, or through Bayesian updating (in which case Bayes' Theorem would give the target function), the subject would go through the following steps: assess a stability distribution over the arguments in Bayes' Theorem; calculate a derived distribution for the posterior from it; and compare that stability distribution with his stability distribution for the direct posterior.




Figure 3-1
STABILITY ADJUSTED RECONCILIATION

Although there would no doubt be bothersome elicitation problems--for example, in eliciting stability dependence between arguments--we would appear to have the basis of a perfectly good procedure for choosing alternative decision analysis or probability models.

In Figure 3-1, if the target were a probability from 0 to 1 and the heights of the lines gave the stability for the direct assessment q and each of two alternative derived readings q' and q'' , the preferred procedure would be q' as the one with highest stability. Note, however, that since the object of this exercise is to choose one approach rather than another and, therefore, which readings to take, the stability distribution must be assessed unconditional on any particular readings since the end product of the exercise is to help decide which readings to take. (This distinction between "pre" and "post" assessments is analogous to that discussed in Brown, 1968 in the context of designing as opposed to interpreting estimates). We might, therefore, distinguish pre from post assessments of reading stability.

Once a target function and its minimally specified readings have been settled upon, the derivation of the target is unique since no reconciliation is needed.

3.3.2 Reconciling an overspecified system after taking readings - What to do, however, in the case of primary interest where overspecified readings have been taken, corresponding, say, to two or more target functions? What do we do when q , q' and q'' have all been elicited? Is there some way to assign a quality measure to all possible target values as represented notionally by the dotted line in Figure 3-1?

It would seem reasonable to equate the quality of the complete reconciled system of assessments with the quality of implied target. (We defer consideration of what to do if there is a target vector rather than a scalar.)

One approach would be to find a measure of tension for any proposed reconciliation and minimize tension. This tension might measure both how incoherent the proposed reconciliation is with all other readings and the stability of the readings. An explicit measure of this tension, however, is still to be developed.

It is possible, as an alternative to tension, that a measure of stability for a reconciliation can be derived which is interpretable in exactly the same way as the stability of an individual reading. If so, then the preferred reconciliation would be the one with the preferred stability distribution, for example, maximum precision. How such a reconciled stability could be determined, however, is not yet clear.

Any particular reconciliation procedure for mapping from raw readings to reconciliation space might itself be interpreted as a function of the readings. For example, one procedure might be to pool different direct and indirect assessments of the same target assessment as a weighted average with weights proportional to the precision of the several target assessments.

In order to achieve reconciliation of all the supporting readings and not just for the target, a procedure would need to be specified for bringing all component readings into conformity with this pooled target value. One which minimizes the summed deviations of the component stability distributions, (measured in standard deviation units) might

be plausible. Any such completely specified reconciliation procedure could be thought of as an analytical function mapping reconciled values onto primary readings and their associated joint stability distribution.

Established probability calculus will determine the distribution of any such function of random variables. Accordingly, the joint stability distribution for any reconciled system of assessment (including target judgment) can then be derived from the distribution of the raw readings and the function corresponding to the reconciliation procedure.

3.4 Use of Measure of Reconciliation Quality

The preferred reconciliation method would seem to be the one which maximizes some measure of total system quality such as stability. (This would not necessarily give unique rationality, but it would give maximum rationality.) No persuasive single scalar measure of stability is apparent to characterize the overall quality of a reconciled system.

However, the reconciled system can be optimized with respect to, say, the stability variance of any one specific target assessment. Therefore, the optimal system and the optimal method of reconciling initial incoherence may depend on which that target assessment is. We may be able to say only that we know in principle how to resolve incoherence for the purpose of a single target assessment, but we may have to acknowledge that a different reconciliation may be appropriate if a different target assessment is involved. This would not allow us to claim any reconciliation universally best, that is, that there is any unique rationality.

If more than one target judgment is to drive the reconciliation, or if it is a vector rather than a scalar (as in the case of a many-valued probability distribution or a multi-attributed utility function), then some more complex measure of the appropriate part of the joint stability distribution needs to be sought for optimization purposes.

A possible approach to reconciled system optimization is to treat the choice of a reconciled assessment system as a decision whose expected opportunity loss is to be minimized, much as one might choose an estimate of probable product demand on which to base a business stocking decision (see Schlaifer 1969). However, the analogy appears brittle when probed. It is not clear how one would define the "loss structure" called for in this type of problem. There appears to be no constructive analogue to the true value with reference to which the loss is to be defined, much less to a probability distribution on divergence from that value.*

* See Notes

4.0 CONCLUSION

In this paper, we have attempted to formulate the scope of a substantially new area of theoretical and applied research and to point to some promising lines of enquiry.

4.1 Bayesian Updating

In principle it appears that we have at least one well-formulated approach--an extension of Bayesian updating where incoherent judgment is treated as data to be processed by higher order judgments--for addressing two of the three key issues discussed in Section 2.1.

Perfect reconciliation is interpreted as the limiting result of a progression of increasingly higher order Bayesian updatings on a data set of readings which is either partial or complete (corresponding to ultimate reconciliation and in some sense perfect rationality).

A practical procedure for partial reconciliation using one level of higher order judgments (prior and likelihood) has been illustrated (in Lindley et al. 1978).

The strategy of seeking out incoherence for reconciliation has not been explicitly addressed, but the general logic for the valuation of differing types and scales of decision analysis (Watson and Brown 1978) appears capable of generalization here.

However, there is a key unresolved theoretical issue of whether and under what circumstances the process of successively higher order judgments converges. Moreover, the practical promise of this approach is limited by the seeming awkwardness of the elicitations called for and by its radical difference from how intelligent subjects in fact appear to resolve incoherence informally.

4.2 Stability-Based Adjustment

An alternative approach based on the stability of initial readings has been discussed but only partially developed. It attempts to model and refine quite closely the intuitively appealing informal reconciliation procedures by which they were initially suggested.

However, the theoretical underpinnings are not clear (nor is it clear how closely it equates or can be reconciled with the Bayesian updating approach). Moreover, no explicit algorithms for achieving reconciliation have yet been proposed--only a principle for choosing among alternative algorithms.

Any process of reconciliation that does not depend solely on the Bayesian paradigm holds some mysteries for us. Does it disperse with some aspect of the Bayesian argument? Or does it restrict the discussion in some way? If so, what could be the nature of either the dispensation or the restriction? Certainly from N's point of view, S and his statements are part of N's external world, and N should presumably process them like any other aspect of his uncertainty. But if N is regarded as, in some way, part of S, then we do have a novel feature not present in the usual formulation of the Bayesian paradigm, namely, an element of introspection that may disturb the situation; though just how is unclear to us. For example, what rules should govern the shiftability? Or how should the different tensions be relaxed?

4.3 What Next?

A great amount of research, theoretical and applied, is immediately indicated, including the following:

- o developing further the conceptual bases discussed here for both Bayesian updating and stability adjustment approaches;
- o developing implementable procedures for a variety of situations under both approaches;
- o testing and developing applied techniques in applied case contexts;
- o investigating the behavioral foundations of incoherence and its reconciliation.

NOTES

(Keyed to Sections of Main Text)

1.1.1 Preparatory steps in reconciling incoherence -It seems reasonable to assume that the subject, expressing incoherent views, has had at least some training in expressing himself probabilistically, so that the grosser errors can be removed. For example, the subject mentioned in the energy example of Section 1.4.2 may be overconfident and unused to expressing the bounds for his judgments, so that the bounds are unrealistically close together. Another may be lacking in confidence and give unusually wide bounds when he is in reality well-informed. The role of training in the removal of some incoherence must not be forgotten. Nor must the role of the psychologist in helping us to understand what types of uncertainty subjects find easy to handle, and what types difficult. All this information is important, and in the Bayesian updating approach described in Section 2.2.3 gets incorporated into N's likelihood for S.

1.2.1 The Personalist Paradigm - Often we shall refer to the personalist paradigm underlying modern decision analysis. By this we mean a view of the world that says that all uncertain situations should be described probabilistically and that probability calculus is therefore the tool for processing uncertainty: some might agree that it is the only tool. In particular, the processing of new information pertaining to an uncertain situation is achieved through Bayes' Theorem. This view of the (uncertain) world will be described as being coherent, so that statements of uncertainty that do not conform to it are incoherent. If decisions are to be included, then an extension of the personalist view admits a utility function, and that decision is to be selected which has the maximum expected utility; the expectation being with respect to the coherent probabilities describing the uncertainty.

2.1 Accurate vs. inaccurate readings on a cognitive field. As a preliminary model of the subject to whom the impeccable analysis is to be inputted, we can posit a large number of potential readings on his judgment of action selection, probabilities, utilities, etc., which represent his cognitive field, typically incoherent.

This may involve problems of interpretation since the measurements are not instantly accessible and the process of measuring them may change the system itself. One can think of two distinguishable types of readings on a target judgment such as a probability: the value as elicited (perhaps mis-measured); and the value correctly elicited (e.g. the subject's true uncertainty), but still possibly incoherent with other correct readings, and subject to reconciliation.

This is the question of whether accuracy of reading should be distinguished from reconciliation of accurate but incoherent readings. Different elicitation techniques can give different readings. Is there a "true" reading (possibly incoherent with other true readings) based on perfect elicitation? For example, one could ask for assessments of uncertainty either as odds or as probabilities and, in general, one would expect different results.

Schlaifer gives a behavioral definition of probability, that is, in terms of the indifference bets and standard lotteries. But doesn't this latter get us back into the problem of assuming rational behavior? If a man bets on Cambridge winning the Boat Race, can we really assume that his probability for Cambridge is higher than 50%? Perhaps there are "higher order effects" that can be ignored.

For the moment we are only concerned with readings as elicited without positing an accurate reading.

2.1.1 Bounding reconciliation space. In principle there is an infinitely large set of reconciled systems, each corresponding to some combination of assessments for the structure in question which are coherent with each other.

In the simplest case, where the structure is $p(A)$ and $1-p(\bar{A})$, any set of complementary probabilities would qualify.

Clearly some bounding on this set is possible. If one has assessed the probability of Cambridge winning as .4 directly and as one minus the probability of Oxford winning or a draw as .7, one would not want to consider reconciled systems of the two assessments which yielded Cambridge win probabilities outside the range .4-.7. However, it is not clear that an acceptable region in reconciliation space should be limited to points derived from raw reading. If there are only two such points, the derived reconciliation should be allowed to lie between the two.

2.1.2 Valuation of decision analysis - It is not now clear with reference to what probability assessments the expectation is taken. If it is to be assessed by the subject, it must somehow relate to his (in principle imperfect) probability assessments. But which of his potentially incoherent probability assessments should he use? The expectation could be taken with respect to the correct probability assessments, but it is not clear what practical value this would have since the subject does not have access to these probabilities.

The expected cost of irrationality, then, gives a value for perfect analysis. Any proposed piece of analysis,

still presumably leading to imperfect results but hopefully less so, would have a value corresponding to the difference between the expected cost of initial imperfect rationality and the expected cost of the new imperfect rationality. However, the new cost of imperfect rationality is a double expectation. The subject expects now what his expected cost of irrationality will be when the proposed analysis is complete. It would appear that the utilities used throughout must be those of perfect analysis. However, since it can be argued that current utility is equal to the expectation of perfect utility, it would appear that either utility can be used interchangeably. (The mathematics of this argument is discussed in Watson and Brown 1978).

2.1.4a Practical procedures for approaching most rational solution. Most applied decision analysts would believe (without necessarily being able to prove) that progressive elaboration of the assessment structure is a good idea. But why? Somehow the idea is that connectivity leads to constraints and therefore stability.

The process of improving on probability assessments includes setting up auxiliary models or functions whose value is the argument of a more primary model. Thus the Oxford/Cambridge assessment might proceed to the conditional assessment conditional on rain, and the probability of rain can then itself be assessed as the output of another indirect model.

Essentially what one is doing is searching for potentially discordant elements in S's external system, that is, that part of the system not yet incorporated into an explicit model. Ideally one would want an assessment which is maximally coherent with the external system. Since an

external system includes elements incoherent with each other, exact coherence between an internal and the external system is not possible. Perhaps something analogous with least-squares fitting would be appropriate, that is, an assessment which does least violence to all other potential measures in the external system, with violence being a function base of distance and of the "validity" of the element it is being confronted with.

It might be that complexity of assessment functions are advantageous because of the greater potential for incoherence, and the more potential incoherence you have the better, but what constitutes "better" is at the heart of our problem. However, it is not clear that the number of arguments in the function is at all the same thing.

2.1.4b Measures of system incoherence. Any particular system not only may display incoherence, but perhaps measurable degrees of it, characterized by something like entropy in engineering systems as has been suggested by Freeman.* This would be a measure of "discordance" in the system. It is not clear that this discordance can be attributed to any part of the system. Intuitively it would seem desirable to seek maximum discordance, say, by increasing the complexity and the overspecification extremes by increasing the number of, say, target functions (but not necessarily the complexity of any particular target function). By analogy with surveying one expects to be better off taking many bearings.

Psychologists have a measure of incoherence called Slater's I, which is used, for example, to measure the degree of incoherence among rankings. A subject is asked

* Peter Freeman, private communication, December 1976

to rank a set of seven objects in terms of, say, probability, but doing it in groups of three, that is, indicating which is the most and which the least probable. When all possible subsets of three have been thus ranked, their implications for the total ranking can be deduced, and in general there will be incoherence. Slater's I gives a measure of this incoherence (See Phillips 1967, 1969; Slater 1960, 1961, 1965).

2.1.4c Choosing a single minimal model. There are no obvious a priori structural grounds for preferring one minimally specified model to another. Complexity is no virtue of itself. Assessing demand for a product as the sum of a large number of additive, say regional components, may be better (whatever that means) than assessing demand directly, but only if in some sense the arguments are more "validly" assessed. One might argue that any direct assessment is a more or less adequate attempt to perform a more indirect disaggregated assessment. By making that process more explicit, one can remove logical errors (the garbage between the garbage in and the garbage out). But this requires direct assessment of the arguments in the function. One could always express one of the arguments as a function of the other arguments and the target value. The notion of veridicality comes in here. For example, one regional market can be assessed or decomposed as the total market less the other regional markets. The natural argument is to say "which arguments does one's experience bear most directly on?" Whatever that may mean.

This is the issue of which single target function to choose. If there is some sense in which one function is preferable to another, then perhaps there is some way of

resolving incoherences between them, which assigns greater weight to the more authoritative model. There is some analogy here to the problem of resolving inconsistencies among probabilistic estimates from different people, for example, weighting them according to the inverse of their variances.

2.2.1a Target functions for an event probability. If the target judgment is a single probability, such as the probability of Cambridge winning the Boat Race, there are a number of different types of target functions, that is, minimally specified structures which imply the target in question.

There is of course direct assessment (though there are different ways of making that assessment, e.g., through betting behavior, odds assessments, probability numbers, etc.). There is a pooling of assessments, e.g., you take a weighted average of different ways of making the elicitation. There is conditioned assessment, e.g., conditioning the Cambridge win probability on rain, or on the results of the toss, or on level of attendance, or anything else, or any combination of these. There is concatenation of target functions, e.g., where the probability of rain required for a single conditioned assessment is itself derived from the quantification of another target function and so forth.

If the target is a many-valued probability distribution then it is a vector, rather than a scalar, as in the case of the probability of the single event, but the basic approaches are the same. On the other hand, if the target is a continuous probability distribution on a scalar, the situation may be a little more complex, unless one equates it to a many-valued discrete distribution (which is

probably realistic, since available measuring instruments have a limit to the fineness with which they can measure, e.g., the nearest cent if it is money).

For continuous distributions there is a further indirect technique, decomposed assessment, where the target scalar about which a distribution is to be assessed can be expressed as an analytical function of two or more arguments, as in the energy usage example given in Section 1.4.2. Strictly speaking, this is not the target function, but the target functions can be deduced from these "decomposition formulas" and a joint distribution on the arguments in them via the calculus of distributions of functions of random variables. Again, target functions can be indefinitely proliferated, for example, by expressing the arguments in one decomposition as a further decomposition themselves. A single example would be demand for a product = number of customers \times average demand per customer. (See Brown et al. 1974, Chapter 34.) Note that the target function may be quite difficult, possibly impractical, to define analytically, but the value can usually be determined to a decent approximation via simulation and other approximating devices. (See Brown 1978.)

Target functions can be interpreted as probability models which are minimally specified. An overspecified model is one where enough inputs are supplied to permit coherence checks. The simplest example of an overspecified model would be one in which both the probability of the event and the probability of its complement are specified, since by coherence one is implied by the other. Similarly, a model substantially more complex may have some parts which

imply others. Such a model can always be re-expressed as two or more target functions, for example, as $p(A)$ or as $1-p(\bar{A})$.

In general, since the subject's elicitation of inputs may not be coherent with each other, the targets derived from two or more target functions will not be the same and there is demonstrable incoherence.

Incoherence may also be generated by the specification of "cross functions," which specify relationships between the arguments in the target functions but do not involve the targets themselves. For example, in the Cambridge win probability case, two target functions might be to express that probability as conditioned assessment with rain and attendance respectively as conditioning events. A cross function might be an assertion of correlation between rain and attendance, which imposes an additional coherence constraint.

The simplest kind of cross function is the complementarity probability of exhaustive events other than the target event.

The sets of all target values implied by all target functions might be described as a feasible target space. Thus there may be no way of formulating questions to the subject which implies a probability of Cambridge winning less than .3 or greater than .9. In this sense, then, we can eliminate some values on the grounds of coherence, and further reconciliation is needed to winnow out the remainder.

There are three distinguishable assessment systems, that is, structures of target and cross functions, with quanti-

fication. The first is the initial set of incoherent readings. The second is a modification of that system to ensure coherence, but without reference to any other part of the subject's external field. The third is the system that emerges from most rational analysis of the entire field. The systems all have the same structure but different values.

2.2.1b Assessment structures and systems. We define as an assessment structure any set of target functions and relevant cross functions that is overspecified, in the sense that once it is quantified it could in principle be demonstrably incoherent. A structure that has been quantified is defined as an assessment system. A raw readings system is one in which the structure had been elicited without regard to coherence from the subject's psycho-field. A reconciled assessment system is one in which coherence has been achieved whether or not it was originally coherent.

2.2.1c Cross-functions. More subtle forms may be constructed which somehow have the function of reducing the freedom of key arguments to slop around. For example, if you start off with two target assessment functions, you might attempt to resolve inconsistencies between them by looking for cross-relationships between their arguments that do not involve the target at all. For example, in the boat race probability example, the two target assessment functions might be conditioned assessment, conditioned respectively on attendance and rain. A third assessment function of a different type might tap the subject's judgment about dependence between the two conditioning events; that is, bad weather is associated with low attendance. The unconditional probabilities of the conditioning events which appear as arguments in the

first two target functions are constrained by the assessment of dependence. However, one still has the problem of what to do when inconsistency is demonstrated.

2.2.2 Unique reconciliation. The question of whether or not there is a unique reconciliation of a set of incoherent statements is clearly related to the broader question of whether a unique statement of uncertainty can be arrived at from any given data set. Even within the personalist paradigm there are two viewpoints: one argues that probability is a type of logical relation between events so that every sensible person will attach the same value to the probability of A given B; the other says that probability is subjective and two coherent observers could differ about this probability.

The logical view is in many ways the more attractive-- and is the one currently popular in statistical treatment of data, though outside the personalist approach. But so far no one has come up with a recipe for how the unique, logical value can be calculated: and this is not despite considerable effort using theories of invariance and other high-powered tools. At the moment we are left with the subjective view, and no unique analysis, so that it seems unlikely that a unique reconciliation is possible with our present knowledge.

Of course, in many situations the conditioning event B is so informative that there is substantial agreement on the value for the probability of A given B, and it seems reasonable to expect that, as we acquire more experience of people as probability assessors, similar practical agreement on the reconciliation procedure will be obtained. What could happen is that N could incorporate this experience into his likelihood for S and hence, with several judgments from S, reach an answer that would not differ by much from that obtained by any other N.

3.2a Assessment vs evidential stability. We would have a quite different measure--for example, variance--depending on whether one is talking about the stability of the psychological assessment or the stability of the evidential base. Thus if you knew you had thought well and hard about a target probability and felt comfortable with it there might be a low pre-posterior variance, that is, high quality for assessment. However, one might simultaneously judge that new evidence would very likely shift the probability substantially and so have low evidence quality.

3.2b Stability dependence. Some thought needs to be given to the interpretation of joint stability assessments. The analogy with joint probability distributions seems quite acceptable, that is, it addresses questions like "If, on further reflection, your assessment of X were to shift in this direction, by this amount, what would happen to your validity distribution on Y?"

3.4 Unique rationality as maximum system stability. If we can in principle derive a measure of validity for an assessed target function, we can presumably also do it for any analytically explicit way of combining target functions (and cross functions) into a reconciled system. This must be so, because the reconciled system is then itself an analytical function of raw assessments.

This is true whether the reconciliation proceeds by a pooling (say, according to least-squares, or a weighting proportional to the reciprocal of validity variances) or by some other reconciliation procedure.

The critical point is: we have an implicit definition of unique rationality if we accept that it corresponds to

the reconciled system with highest validity. We only need to be able to specify all possible reconciliation procedures to determine that which maximizes system validity.

We say "only," though the mind boggles at the practical difficulties of implementing such a procedure. However, in this paper we are only concerned to identify theoretical principles. (In particular the practical handling of validity dependences would be horrendous.)

We are, however, still left with the philosophical problem of defining a stability distribution in such a way that it does not assume the rationality reconciliation procedure it is being used to define.

If the shift in assessment is predicated on "perfect rationality," can we use it to define perfect rationality? Possibly we can in some iterative, convergent, asymptotic way.

GLOSSARY

(Underlined terms in explanation are explained elsewhere in Glossary)

Adjusted Reading - reading adjusted, e.g., to achieve coherence.

Assessment - a value (for probability, utility, choice) applied by a subject to an object.

Assessment Structure - a group of related target functions (and cross-functions), a model without specified assessments but with potential incoherence.

Assessment System - quantified structure (i.e., with specific assessments).

Bayesian Updating - use of Bayes' Theorem to derive posterior from prior and likelihood.

Coherence - logical compatibility (e.g., according to probability calculus).

Decomposition - expressing a variable as an analytic function of other variables (e.g., demand per customer x demand per customer).

Elicitation - taking a reading on an element in an assessment system (e.g., probabilities, utilities).

First Order Readings - quantities of direct interest.

Minimal Assessment Structure - one with no potentiality for incoherence.

Normative Investigator (N) - the source of second order readings elicited to assure coherence in first order readings--a partition of the subject's cognitive field.

Object - a real-world entity such as event, act, relationship.

Optimal System/Target - "most rational" target assessment (and embedding system).

Partial Reconciliation ($\hat{\pi}$) - an attempt at estimating perfect reconciliation π , based on subset (q) of all potential readings (Q).

Perfect Analysis or Reconciliation (π) - the result of applying unique rationality to all potential readings (Q).

Precision - second order measure of the validity of a reading.

Psychological Field (or Psycho-Field) - everything in S's head--totality of actual or potential readings available for elicitation.

Reading or Raw Reading (Q) - a number (e.g., probability) elicited straight from S's field, that is, unconstrained by coherence.

Reconciled Assessment System - any coherent reconciliation of incoherent raw readings.

Second Order Readings (P) - readings taken to assure coherence in first order readings, themselves adjusted to be coherent.

Stability - a measure of the validity of an assessment (e.g., probability distribution of shift in assessment on further analysis).

Stability Adjustment - reconciliation method based on stability of raw readings.

Subject (S) - the person whose judgments are analyzed (at a given point in time unless otherwise specified).

Target (T) - an object or assessment the subject is primarily interested in (e.g., a posterior probability).

Target Function - algorithm (e.g., Bayes' Theorem) deriving target (e.g., posterior) from other assessments (e.g., prior, likelihoods).

Target Space - set of possible target assessments.

Ultimate Reconciliation (π) - perfect reconciliation of all a subject's potential readings.

Unique Rationality - the concept that subject has a single most coherent interpretation of his total cognitive field.

Validity - measure of the quality (e.g., stability) of reading (raw or derived).

REFERENCES

- Brown, R. V. "Credence Analysis: a Guide to Choosing Research Strategies with Particular Reference to Industrial Market Research." Metra, Volume 3, 1963.
- Brown, R. V. "Heresy in Decision Analysis: Modelling Subsequent Acts Without Rollback." Decision Sciences. October 1978 (in press).
- Brown, R. V. Research and the Credibility of Estimates. Harvard University, Graduate School of Business Administration, Division of Research. Boston, Massachusetts: 1969.
- Brown, R. V.; Kahr, A. S.; and Peterson, C. R. Decision Analysis for the Manager. New York: Holt, Rinehart and Winston, 1974.
- Brown, R. V., and Lindley, D. V. "Notes on Rationality and the Reconciliation of Incoherence" in Watson, S. R.; Brown, R. V.; and Lindley, D. V. Three Papers on the Valuation of Decision Analysis. Technical Report 77-2. McLean, Virginia: Decisions and Designs, Inc. 1977. (Office of Naval Research Decision Research Programs).
- Brown, R. V., and Ulvila, J. W. Selecting Analytic Approaches for Decision Situations (Revised Edition). Technical Report TR 77-7-25. McLean, Virginia: Decisions and Designs, Inc., December 1977.
- French, S. Updating of Belief in the Light of Someone Else's Opinion. Department of Decision Theory, University of Manchester, April 1978, (submitted to the Journal of the Royal Statistical Society).
- Lindley, D. V. Making Decisions. New York: Wiley-Interscience, 1973.
- Lindley, D. V.; Tversky, A.; and Brown, R. V. On the Reconciliation of Probability Assessments. Technical Report 77-10-25. McLean, Virginia: Decisions and Designs, Inc., February 1978. (accepted by the Journal of the Royal Statistical Society for publication in 1979).

- Morris, P. A. "Decision Analysis Expert Use." Management Science, Volume 20, Number 9, May 1974.
- Novick, M. R., and Lindley, D. V. "The Use of More Realistic Utility Functions in Educational Applications." Journal of Educational Measurement (in press), 1978.
- Novick, M. R., and Lindley, D. V. "Fixed State Utility Assessments". Journal of the American Statistical Association (in preparation).
- Phillips, J. P. M. British Journal of Mathematics & Statistical Vol. 20, 1967, pp 217-225.
- Phillips, J. P. M. British Journal of Mathematics & Statistical Psychology. Vol. 22, 1969, pp 97-101.
- Raiffa, H., and Schlaifer, R. C. Applied Statistical Decision Theory. Harvard University, 1961.
- Savage, L. J. Foundations of Statistics. New York: Dover Publications, 1972.
- Schlaifer, R. O. Analysis of Decision Under Uncertainty. New York. McGraw-Hill, 1969.
- Slater, P. Brit. J. Statis. Psychol. XIII, 1960, 119-135.
- Slater, P. Biometrika. 48, 1969, 303-312.
- Slater, P. Brit. J. Math. & Statis. Psychol. Vol. 18, 1965, 227-242.
- Sturrock, P. A. Astrophysical Journal. Vol. 182, 1973 pp 569-580.
- Watson, S. R., and Brown, R. V. "The Valuation of Decision Analysis." Journal of the Royal Statistical Society. Series A, 1978 (in press).
- Watson, S. R.; Weiss, J. J.; and Donnell, M. L. "Fuzzy Decision Analysis." McLean, Virginia: Decisions and Designs, Inc., May 1978. (Submitted to the IEEE Transactions on Systems, Man, and Cybernetics.)
- Zadeh, L. H. "Theory of Fuzzy Sets." Memorandum No. UCB/ERL M77-1. University of California, Berkeley, California. January 1977.

CONTRACT DISTRIBUTION LIST
(Unclassified Technical Reports)

Director Advanced Research Projects Agency Attention: Program Management Office 1400 Wilson Boulevard Arlington, Virginia 22209	2 copies
Office of Naval Research Attention: Code 455 800 North Quincy Street Arlington, Virginia 22217	3 copies
Defense Documentation Center Attention: DDC-TC Cameron Station Alexandria, Virginia 22314	12 copies
DCASMA Baltimore Office Attention: Mr. K. Gerasim 300 East Joppa Road Towson, Maryland 21204	1 copy
Director Naval Research Laboratory Attention: Code 2627 Washington, D.C. 20375	6 copies

SUPPLEMENTAL DISTRIBUTION LIST
(Unclassified Technical Reports)

Department of Defense

Director, Defense Advanced Research
Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Director, Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Dr. Don Hirta
Defense Systems Management School
Building 202
Ft. Belvoir, VA 22060

Department of the Navy

Office of the Chief of Naval Operations
(OP-951)
Washington, Dc 20450

Office of Naval Research
Assistant Chief for Technology (Code 200)
800 N. Quincy Street
Arlington, VA 22217

Office of Naval Research (Code 230)
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
Naval Analysis Programs (Code 431)
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
Operations Research Programs (Code 434)
800 North Quincy Street
Arlington, VA 22217

Office of Naval Research
Information Systems Program (Code 437)
800 North Quincy Street
Arlington, VA 22217

Director, ONR Branch Office
Attention: Dr. Charles Davis
536 South Clark Street
Chicago, IL 60605

Director, ONR Branch Office
Attention: Dr. J. Lester
495 Summer Street
Boston, MA 02210

Director, ONR Branch Office
Attention: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Director, ONR Branch Office
Attention: Mr. R. Lawson
1030 East Green Street
Pasadena, CA 91106

Office of Naval Research
Scientific Liaison Group
Attention: Dr. M. Bertin
American Embassy - Room A-407
APO San Francisco 96503

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps (Code RD-1)
Washington, DC 20380

Naval Personnel Research and Development
Center (Code 305)
Attention: LCDR O'Bar
San Diego, CA 92152

Department of the Army

U. S. Army Research Institute
Organizations and Systems Research
Laboratory
Attention: Dr. Edgar M. Johnson
5001 Eisenhower Avenue
Alexandria, VA 22333

Department of the Air Force

Air Force Office of Scientific Research
Life Sciences Directorate
Building 410, Bolling AFB
Washington, D. C. 20032

Other Institutions

Director, Social Science Research
Institute
University of Southern California
Attention: Dr. Ward Edwards
Los Angeles, CA 90007

Perceptronics, Incorporated
Attention: Dr. Amos Freedy
6271 Variel Avenue
Woodland Hills, CA 91364

Stanford University
Attention: Dr. R. A. Howard
Stanford, CA 94305

Director, Applied Psychology Unit
Medical Research Council
Attention: Dr. A. D. Baddeley
15 Chaucer Road
Cambridge, CB 2EF
England

Department of Psychology
Brunel University
Attention: Dr. Lawrence D. Phillips
Uxbridge, Middlesex UB8 3PH
England

Decision Analysis Group
Stanford Research Institute
Attention: Dr. Miley W. Merkhofer
Menlo Park, CA 94025

Decision Research
1201 Oak Street
Eugene, OR 97401

Department of Psychology
University of Washington
Attention: Dr. Lee Roy Beach
Seattle, WA 98195

Department of Electrical and Computer
Engineering

University of Michigan
Attention: Professor Kan Chen
Ann Arbor, MI 94135

Department of Psychology
Hebrew University
Attention: Dr. Amos Tversky
Jerusalem, Israel

Professor Howard Raiffa
Morgan 302
Harvard Business School
Harvard University
Cambridge, MA 02163

Department of Psychology
University of Oklahoma
Attention: Dr. Charles Gettys
455 West Lindsey
Dale Hall Tower
Norman, OK 73069

Institute of Behavioral Science #3
University of Colorado
Attention: Dr. Kenneth Hammond
Room 201
Boulder, Colorado 80309

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER TR 78-8-72 ✓	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) RECONCILING INCOHERENT JUDGMENTS (RIJ) --TOWARD PRINCIPLES OF PERSONAL RATIONALITY		5. TYPE OF REPORT & PERIOD COVERED Technical
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Rex V. Brown Dennis V. Lindley, University College, London		8. CONTRACT OR GRANT NUMBER(s) Contracts N00014-75-C-0426 and N00014-78-C-0100
9. PERFORMING ORGANIZATION NAME AND ADDRESS Decisions and Designs, Incorporated 8400 Westpark Drive, Suite 600, Box 907 McLean, Virginia 22101		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		12. REPORT DATE July 1978
		13. NUMBER OF PAGES 82
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research 800 North Quincy Street Arlington, Virginia 22217		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision theory, cognitive psychology, philosophy		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Decision analysis involves constructing models which force logical coherence between a subject's judgments, e.g., between his choice of action and probabilities and utilities. However, it does not specify how he should reconcile any incoherent judgments. There are indefinitely many ways they can be adjusted to be coherent systems of judgment. The authors discuss two approaches for identifying one ideal set of reconciled judgments for a subject, given some or, in the limit, all potential incoherent "readings." They both call for higher order judgments bearing on the "precision" of the subject's original readings.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

One is a straightforward extension of Bayesian updating with the readings serving as data to update a prior. The other involves minimizing adjustments, taking into account the stability of the readings as probabilistically measured.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)